

What is claimed is:

3 1. In a data packet transmission system, a method of filtering an incoming data packet
4 comprising the steps of:

5 creating a decision graph having a tree-like hierarchy for a set of filters such that
6 the filters are subdivided and grouped along edges and stages of the tree-like hierarchy
7 according to common filter values;

8 receiving a data packet at a point of ingress to the data packet transmission system;

9 comparing filters along edges of the tree-like hierarchy to portions of the data
10 packet to determine where there are matches; and

11 if matches are made along an entire edge of the tree-like hierarchy, filtering the data
12 packet stream according to number of condition values compared.

13 2. The method according to claim 1, wherein said creating step is performed so that
14 the tree-like hierarchy includes wildcards at specified byte values.

15 3. The method according to claim 2, wherein said creating step is performed by the
16 step of subdividing the set of filters according to common values at specified bytes of said
17 filters.

18 4. The method according to claim 3, wherein the comparing step further comprises the
19 step of dynamically choosing which byte of a filter to compare against portions of the
20 classification key.

21 5. The method according to claim 4, further comprising the step of choosing a next
22 byte to compare based on the number of wildcards in a filter.

1 6. The method according to claim 5, further comprising the step of choosing the byte
2 with the most wildcards.

1 7. The method according to claim 1, further comprising the step of classifying the data
2 packet according to the number of condition value matches encountered during
3 comparison with the tree-like hierarchy.

1 8. The method according to claim 1, wherein the condition values are limited by a
2 number of bytes in a filter within the groups of filters.

1 9. The method according to claim 1, wherein a size of the tree-like hierarchy is
2 minimized based on a specified criteria.

1 10. The method according to claim 9, wherein the specified criteria is based on a
2 number of irrelevant condition values in the data packet stream.

1 11. The method according to claim 1, wherein the step of creating a decision tree is
2 performed such that the set of filters are divided into groups of filters based on specific
3 bytes in a search key.

1 12. The method according to claim 11, wherein the set of filters are divided based on
2 comparing one byte at a time in the search key.

1 13. The method according to claim 1, wherein the filter conditions are divided into
2 groups of filters based on a maximum number of bits which all filters in the groups of filters
3 have in common.

- 1 14. The method according to claim 1, further comprising:
2 storing intermediate matching values; and
3 matching the stored results with a next byte of the data packet.
- 1 15. In connection with a router that receives incoming packets from a data network, a
2 method of filtering the incoming packets at the interface of the router to provide
3 differentiated services, the method comprising the steps of:
4 constructing a decision tree structure that stores a set of filters in a tree-like
5 hierarchy, the tree-like hierarchy including edges and stages, the set of filters subdivided
6 according to their values along the edges and stages of the tree-like hierarchy;
7 the interface of the router receiving at least one incoming data packet; and
8 comparing segments of the incoming data packet to values of the filters stored in
9 the tree-like hierarchy such that the segments are compared along stages and not every
10 edge of the tree-like hierarchy to allow the incoming data packet to be differentiated
11 according to matches between the segments compared and specific filter values.
- 1 16. The method of claim 15 wherein said comparison step further comprises the step
2 of executing a single instruction within the PPU that filters the incoming data packet
3 according to its content.
- 1 17. The method of claim 16 wherein the step of executing the single instruction further
2 comprises the step of advancing from the beginning of the incoming data packet by a
3 predetermined amount of offset before beginning to compare.

1 18. The method of claim 15 wherein said comparing step comprises the further steps
2 of:

3 comparing a first specific byte of the incoming data packet to values of filters in the
4 upper most stage of the tree-like hierarchy;

5 once a match is made, comparing a second specific byte of the incoming data
6 packet to values of filters contained in a stage below the upper most stage of the tree-like
7 hierarchy and along edges that connect to filters in the upper most stage that match in
8 values to the first specific byte.

1 19. The method of claim 18 further comprising the step of comparing further bytes of
2 the incoming data packets to values of filters contained in subsequent stages of the tree-
3 like structure until a match is made at the lowest stage of the tree-like structure or no
4 match at all.

1 20. The method of claim 19 wherein said comparing steps further comprise the step of
2 dynamically choosing the next byte of the incoming data packet to compare.

1 21. The method of claim 19 wherein said comparing steps further comprise the step of
2 replacing specified values in the filters with wildcards.

1 22. The method of claim 19 wherein said comparing steps further comprise the step of
2 collapsing values in any stage of the tree-like structure so that multiple branches from a
parent can lead to the same child.

- 1 23. A router for use in a data network and for allowing the implementation of
2 differentiated services by filtering incoming data packets, the router comprising:
3 an interface to the data network;
4 a packet processing unit coupled to said interface and adapted for receiving
5 incoming data packets via the interface from the data network;
6 memory means operably coupled to said packet processing unit and adapted for
7 storing software instructions that cause said packet processing unit to filter incoming data
8 packets; and
9 a data structure for storing a set of filters in a tree-like hierarchy, the tree-like
10 hierarchy including edges organized into stages, the set of filters subdivided according to
11 their values along the edges and stages of the tree-like hierarchy;
12 wherein said software instructions are further adapted to cause said packet
13 processing unit to compare segments of incoming data packets to values of the filters
14 stored in the tree-like hierarchy such that the segments are compared along stages and
15 not every edge of the tree-like hierarchy to allow the incoming data packet to be
16 differentiated according to matches between the segments compared and specific filter
17 values.
- 1 24. The router of claim 23 wherein said packet processing unit is an application specific
2 integrated circuit or general purpose CPU.
- 1 25. The router of claim 24 wherein said application specific integrated circuit is further
2 adapted to filter incoming IP packets based on the contents of the IP packets and using
3 said tree-like structure to provide differentiated services.
- 1 26. The router of claim 23 wherein said software instructions comprise a single
2 instruction capable of causing said packet processing unit to filter incoming data packets.

1 27. The router of claim 23 wherein said software instructions are further adapted to
2 cause said packet processing unit to compare specific bytes of incoming data packets with
filter values of said tree-like structure.

1 28. The router of claim 27 wherein said software instructions are further adapted to
2 cause said packet processing unit to compare a first specific byte of the incoming data
3 packet to values of filters in the upper most stage of the tree-like hierarchy and once a
4 match is made to compare a second specific byte of the incoming data packet to values
5 of filters contained in a stage below the upper most stage of the tree-like hierarchy and
6 along edges that connect to filters in the upper most stage that match in values to the first
7 specific byte.